

CPSC 311: Definition of Programming Languages: Static semantics: Typing rules ("13-typing-rules")

Joshua Dunfield
University of British Columbia

October 18, 2015

Code

The code for the first part of these lecture notes can be found in

<http://www.ugrad.cs.ubc.ca/~cs311/2015W1/notes/typing-with.rkt>

Code for the last part (Section 5) can be found in

<http://www.ugrad.cs.ubc.ca/~cs311/2015W1/notes/typing-lam.rkt>

Topics discussed

- when type checking happens
- does an ill-typed program mean anything?
- typing judgments; typing contexts
- typing rules for AE, WAE, BWAE
- typing for lam
- modes: which meta-variables are inputs, and which are outputs?
- declarative vs. algorithmic
- one solution: require types inside lam and rec

1 When does typing happen?

Typing is part of *static* semantics, and evaluation is part of *dynamic* semantics, so the type checker has to run before the interpreter. In most language implementations, type checking is done after parsing. We can view the parser and type checker as increasingly restrictive filters: the parser

§1 When does typing happen?

only accepts programs that match a BNF, and the type checker only accepts programs that are well-typed (that can be typed according to a system of rules).

Some “advanced” type systems are implemented not as replacements of the type checker, but as additional type checkers that are run after the “normal” one. For example, the SML-CIDRE “sort checker” is run after the usual SML type checker, providing a *third* filter on what counts as a valid program.

2 Are ill-typed programs meaningful?

Should we regard a program that isn’t well-typed as having any meaning at all?

Taking the definition of a language to be “syntax + static semantics + dynamic semantics”, then a reasonable answer is “no”: Just as we don’t try to give any kind of semantics (static or dynamic) to strings that are not accepted by the parser, we shouldn’t try to give a dynamic meaning to programs that are not accepted by the type checker.

However, we could interpret the question as, “If I imagine a *different* language that’s the same but with no type system, does the program have a meaning?” In that case, there is no type system to filter out programs, and we can certainly ask what the *dynamic* meaning of the program is; we’ve been doing this for [Untyped] Fun for weeks now!

A tricky but more interesting version of the question arises in languages that can have *more than one type* for each expression, that is, when $\Gamma \vdash e : A$ and $\Gamma \vdash e : B$ and $A \neq B$. For example, in Java, every object is a subclass of `Object`, and so a variable of type `Integer` has both the type `Integer` and the type `Object`. The original question could be rephrased as, “Does an ill-typed program have zero or one meanings?” Now the question becomes, “Does a program with more than one type have more than one meaning?” The answer is a matter of taste. We could say e has *two* meanings, A and B , or that the (static) meaning of e is not one type, but the set of types $\{A, B\}$. Whether we choose to say that e has two static meanings or one, we still need to relate the static meaning(s) of e to the dynamic meaning of e (what it evaluates to).

For now, all of our type systems will have the property that each expression has at most one type. But you should keep in mind that this is a property of these particular type systems, not of all type systems.

In the case of Java, we could have only one static meaning by saying that the meaning of an expression is its smallest type (that is, its “lowest” class). This is sometimes called the *principal type* of the expression.

3 Defining a type system

We didn’t specify any static semantics for Fun++, so Fun++ is untyped. It’s time to design *Typed* Fun++.

This will enable us to catch many errors statically (after parsing) rather than during evaluation. It will also let us write type annotations (signatures) in Fun++ programs that, unlike signatures

§3 Defining a type system

in comments in Racket, are actually checked.

Again, a language is syntax *and* semantics. Typed Fun++ will have almost the same syntax as Fun++; the only change will be (in the abstract syntax) variants for type annotations. The important difference will be in the static semantics.

When we make our language typed, does it affect the dynamic semantics? We'll find out!

3.1 Typing judgment

The usual judgment form for whether an expression is typed, and which type it has, is

$\boxed{\Gamma \vdash e : A}$ Under assumptions Γ , expression e has type A

The symbol \vdash is called a *turnstile*; it separates the judgment into assumptions on the left, and something that—if the whole judgment is derivable—is a logical consequence of the assumptions. (Thus, at a very high level, it means “implies”; but that does not distinguish it from many other notions in logic and programming languages, including Gentzen’s horizontal lines.)

Ignoring the Γ for the moment, the judgment says that the given expression e has the type A . In the dynamic semantics of Fun++, the operators (+, =, etc.) evaluate only if they are applied to appropriate kinds of animals (numbers vs. booleans vs. functions), so our type system should, at minimum, distinguish numbers and booleans. We would then expect

$$\Gamma \vdash (\text{num } 3) : \text{Num}$$

and

$$\Gamma \vdash (\text{bfalse}) : \text{Bool}$$

to be derivable.

Just these two types Num and Bool would suffice for a language without functions. But to handle functions, or even just with, the judgment form needs to talk about the types of the expression’s free identifiers. Our dynamic semantics (big-step $e \Downarrow v$ and small-step $e_1 \longrightarrow e_2$) didn’t need assumptions, because those rules substituted away identifiers. So, assuming we start with a program e that is closed (has no free identifiers), we never need to evaluate (or step) any expression with free identifiers.

On the other hand, types are *static* semantics—meanings given to expressions without evaluating them. Given a function (lam x e), we can’t substitute a value for x , because we don’t know what value to use until the function is applied, and we don’t know how the function will be applied without evaluating the whole program.

To give a *static* meaning—a type—to a function body, or to the body of a with, we need to handle expressions that have free identifiers. Such expressions are called *open*, because they are not closed.

The assumptions Γ , also called a *typing context*, simply list the types of the free identifiers. As we recently did for values and evaluation contexts (for small-step semantics), we can use a BNF grammar to define what a Γ is. As with values and evaluation contexts, we are *not* specifying

§3 Defining a type system

concrete syntax. In fact, we’re going one step further away from using a BNF for concrete syntax, because Γ has no direct connection to the syntax of the language. Rather, we are using the BNF notation for a different purpose.

Typing contexts $\Gamma ::= \emptyset$ empty context (no assumptions)
 $| x : A, \Gamma$ x has type A , with more assumptions

This BNF defines a list of assumptions: you can think of \emptyset as the empty list, and $x : A, \Gamma$ as the “cons” of a head $x : A$ and a tail Γ . Given this BNF, we should not be able to write

$x : A$

alone for a typing context with a single assumption; rather, we should write

$x : A, \emptyset$

By convention, however, typing contexts are treated somewhat more informally, so that $x : A$ and $x : A, y : B$ would be considered valid contexts. However, we will follow the BNF exactly when we encode it as a **define-type**. This BNF represents abstract syntax, not as concrete syntax, so we encode it using **define-type**. We don’t need to write a parser for Γ because Γ won’t appear in programs, only in the typing judgment.

3.2 Typing for AE

To design the typing rules, we can “replay” most of the development that led us to `Fun++`: start with `num`, `add` and `sub`; `add with`; `add functions`; `add recursion`; `add booleans` and `ite`; `add pairs`.

Until we reach `with`, we won’t actually need Γ —it will always be empty—but I’m adding it now to smooth the transition to the larger language.

The rules for AE expressions are given in Figure 1. There are only three rules... and only one type! Having only one type isn’t very useful. (These rules are, in a sense, only repeating the structure of the AE abstract syntax itself; if you read “: Num” as “is an AE”, the rules become a less compact notation for **define-type**.) When we put booleans back, we’ll add a second type `Bool`, and then the types will really tell us something. In fact, the types will matter earlier than that...

$\Gamma \vdash e : \text{Num}$ Under assumptions Γ (always empty, here), AE expression e has type `Num`

$$\frac{}{\Gamma \vdash (\text{num } n) : \text{Num}} \text{AE-Type-num}$$
$$\frac{\Gamma \vdash e1 : \text{Num} \quad \Gamma \vdash e2 : \text{Num}}{\Gamma \vdash (\text{add } e1 \ e2) : \text{Num}} \text{AE-Type-add} \quad \frac{\Gamma \vdash e1 : \text{Num} \quad \Gamma \vdash e2 : \text{Num}}{\Gamma \vdash (\text{sub } e1 \ e2) : \text{Num}} \text{AE-Type-sub}$$

Figure 1 Typing rules for AE expressions

§3 Defining a type system

3.3 Typing for WAE

Adding with and id to the AE language gives us the WAE language. We can copy the three rules for AE expressions, adding a W in front of the rule names. But now Γ will serve a purpose. We type the body of a with using a new assumption $x : \text{Num}$, and when we type $(\text{id } x)$, we check that “ $(x : \text{Num}) \in \Gamma$ ”, that is, that $x : \text{Num}$ appears somewhere in Γ .

It would be possible, but tedious, to say that $(x : \text{Num}) \in \Gamma$ is another form of judgment, and write rules deriving it.

$\boxed{\Gamma \vdash e : \text{Num}}$ Under assumptions Γ (not always empty, now!), WAE expression e has type Num

$$\begin{array}{c} \frac{}{\Gamma \vdash (\text{num } n) : \text{Num}} \text{WAE-Type-num} \\ \\ \frac{\Gamma \vdash e1 : \text{Num} \quad \Gamma \vdash e2 : \text{Num}}{\Gamma \vdash (\text{add } e1 \ e2) : \text{Num}} \text{WAE-Type-add} \quad \frac{\Gamma \vdash e1 : \text{Num} \quad \Gamma \vdash e2 : \text{Num}}{\Gamma \vdash (\text{sub } e1 \ e2) : \text{Num}} \text{WAE-Type-sub} \\ \\ \frac{\Gamma \vdash e : \text{Num} \quad x : \text{Num}, \Gamma \vdash e\text{Body} : \text{Num}}{\Gamma \vdash (\text{with } x \ e \ e\text{Body}) : \text{Num}} \text{WAE-Type-with} \quad \frac{(x : \text{Num}) \in \Gamma}{\Gamma \vdash (\text{id } x) : \text{Num}} \text{WAE-Type-var} \end{array}$$

Figure 2 Typing rules for WAE expressions

We still only have one type, but our typing rules are no longer useless: they are checking that all free variables in the expression appear in Γ . When we begin typing an expression, we are not inside any with expression, so Γ is \emptyset ; as we enter the body of a with, we add the variable now in scope.

Consequently, if we *can* derive $\emptyset \vdash e : \text{Num}$, then evaluating e *cannot* raise a free-variable-error. For example, evaluating $(\text{with } x \ (\text{id } y) \ (\text{id } x))$ will raise a free-variable-error because y is free, but that expression is rejected by typing: we cannot derive

$$\emptyset \vdash (\text{with } x \ (\text{id } y) \ (\text{id } x)) : \text{Num}$$

■ **Exercise 1.** Try to derive the above judgment.

■ **Exercise 2.** Derive $y : \text{Num} \vdash (\text{with } x \ (\text{id } y) \ (\text{id } x)) : \text{Num}$.

■ **Exercise 3.** Derive $\emptyset \vdash (\text{with } y \ (\text{num } 2) \ (\text{with } x \ (\text{id } y) \ (\text{id } x))) : \text{Num}$.

(If you derived a judgment in a previous exercise, and that judgment appears as a premise, just write a checkmark above the premise.)

Since we only have one type, we have no “agreement errors”, but the type system does catch the one kind of error we have in WAE!

§3 Defining a type system

3.4 WAE + booleans

To add booleans (bfalse, btrue, ite), we need more than one type, so it's time to give a BNF grammar for types as well. Whether this BNF grammar is serving as concrete syntax (like the grammar for $\langle E \rangle$) or another notation for **define-type** can be set aside for now, because the grammar is so simple; we'll have to revisit this later, much to my annoyance.

$$\begin{array}{l} \text{Types } A, B ::= \text{Num } \text{numbers} \\ \quad \quad \quad | \text{Bool } \text{booleans} \end{array}$$

Compared to WAE, some of the rules don't change at all, and some are completely new (for bfalse, btrue, ite). BWAE-Type-with and BWAE-Type-var have the same structure, but instead of the single type Num everywhere, these rules have meta-variables A and B, allowing with to bind a Num in a body of type Bool, or a Bool in a body of type Num, or any other combination.

Writing A and B in BWAE-Type-with doesn't mean that A and B are necessarily different types, only that they don't have to be the same type. If we wanted to require A and B to be different, we would need to add a premise $A \neq B$.

$\boxed{\Gamma \vdash e : A}$ Under assumptions Γ , WAE+booleans expression e has type A

$$\begin{array}{c} \frac{}{\Gamma \vdash (\text{num } n) : \text{Num}} \text{BWAE-Type-num} \\ \\ \frac{\Gamma \vdash e1 : \text{Num} \quad \Gamma \vdash e2 : \text{Num}}{\Gamma \vdash (\text{add } e1 \ e2) : \text{Num}} \text{BWAE-Type-add} \quad \frac{\Gamma \vdash e1 : \text{Num} \quad \Gamma \vdash e2 : \text{Num}}{\Gamma \vdash (\text{sub } e1 \ e2) : \text{Num}} \text{BWAE-Type-sub} \\ \\ \frac{}{\Gamma \vdash (\text{bfalse}) : \text{Bool}} \text{BWAE-Type-false} \quad \frac{}{\Gamma \vdash (\text{btrue}) : \text{Bool}} \text{BWAE-Type-true} \\ \\ \frac{\Gamma \vdash e : \text{Bool} \quad \Gamma \vdash e\text{Then} : A \quad \Gamma \vdash e\text{Else} : A}{\Gamma \vdash (\text{ite } e \ e\text{Then} \ e\text{Else}) : A} \text{BWAE-Type-ite} \\ \\ \frac{\Gamma \vdash e : A \quad x : A, \Gamma \vdash e\text{Body} : B}{\Gamma \vdash (\text{with } x \ e \ e\text{Body}) : B} \text{BWAE-Type-with} \quad \frac{(x : A) \in \Gamma}{\Gamma \vdash (\text{id } x) : A} \text{BWAE-Type-var} \end{array}$$

Figure 3 Typing rules for BWAE expressions

■ **Exercise 4.** In assignment 2, we replaced add and sub with binop and then added lessthanop and equalsop. Suppose we decide not to use binop, but instead add an abstract syntax variant lessthan that behaves like a2's (binop (lessthanop) ...). Design a rule "BWAE-Type-lessthan" that types expressions of the form (lessthan $e1 \ e2$).

§3 Defining a type system

3.5 All the typing rules (that we can't implement)

Adding rules for functions (Type-lam and Type-app) below is pretty straightforward on paper. Unfortunately, we can't implement Type-lam! The problem is that our type checker is only given Γ and e . We know that e has the form $(\text{lam } x \ e\text{Body})$, so we know that we need to apply rule Type-lam, but we don't know what A is, so we don't know what we need to derive next!

This leads us to the concept of the *mode* of a meta-variable.

(Aside: In the rule Type-binop, I am assuming a judgment form $\text{op} : A1 * A2 \rightarrow B$, read “operator op takes two arguments of types $A1$ and $A2$, respectively, and returns a value of type B ”.)

$\Gamma \vdash e : A$ Under assumptions Γ , expression e has type A

$$\begin{array}{c} \frac{}{\Gamma \vdash (\text{num } n) : \text{Num}} \text{Type-num} \quad \frac{(x : A) \in \Gamma}{\Gamma \vdash x : A} \text{Type-var} \\ \frac{\text{op} : A1 * A2 \rightarrow B \quad \Gamma \vdash e1 : A1 \quad \Gamma \vdash e2 : A2}{\Gamma \vdash (\text{binop } \text{op } e1 \ e2) : B} \text{Type-binop} \\ \frac{}{\Gamma \vdash (\text{bfalse}) : \text{Bool}} \text{Type-false} \quad \frac{}{\Gamma \vdash (\text{btrue}) : \text{Bool}} \text{Type-true} \\ \frac{\Gamma \vdash e : \text{Bool} \quad \Gamma \vdash e\text{Then} : A \quad \Gamma \vdash e\text{Else} : A}{\Gamma \vdash (\text{ite } e \ e\text{Then} \ e\text{Else}) : A} \text{Type-ite} \\ \frac{x : A, \Gamma \vdash e\text{Body} : B}{\Gamma \vdash (\text{lam } x \ e\text{Body}) : A \rightarrow B} \text{Type-lam} \quad \frac{\Gamma \vdash e1 : A \rightarrow B \quad \Gamma \vdash e2 : A}{\Gamma \vdash (\text{app } e1 \ e2) : B} \text{Type-app} \\ \frac{\Gamma \vdash e : A \quad x : A, \Gamma \vdash e\text{Body} : B}{\Gamma \vdash (\text{with } x \ e \ e\text{Body}) : B} \text{Type-with} \end{array}$$

Figure 4 Typing rules for Typed Fun++... not implementable!

3.5.1 One judgment, many problems

Each judgment form is defined by the rules that can derive it, that is, the rules that have that judgment form as the rule's conclusion. Both philosophically and practically, each judgment form exists independently from its *implementation* as an evaluator, or stepper, or type checker. A more subtle point is that each judgment has several different implementations that solve totally different problems!

For example, our interpreter implements the evaluation judgment $e \Downarrow v$. But the problem our interpreter solves—or, equivalently, the question that running our interpreter answers—is really

§3 Defining a type system

“Given an expression e , is there some value v such that $e \Downarrow v$?” (Interpreter question)

We could ask a different question about that judgment:

“Given a value v , is there some expression e such that $e \Downarrow v$?” (Reverse interpreter question)

This is an easier question—it can be answered using a single line of Racket code. However, a similar-looking question is much harder to answer than the interpreter question!

“Given a value v , list *all* the expressions e such that $e \Downarrow v$.” (Horrible question)

(I’m not sure that question is even decidable. . . there are certainly too many such expressions to have any hope of listing them all!)

We could also ask

“Given an expression e and a value v , is the judgment $e \Downarrow v$ derivable?” (Validation question)

Here, the instantiations of both meta-variables are given.

At the other extreme, we could ask

“Do there exist e and v such that $e \Downarrow v$ is derivable?” (Vacuousness question)

Answering this question is easy: pick a rule with no premises, like Eval-num, and choose any n you like. (But it’s not a completely pointless question. Mathematicians tell a story about a PhD student who proved many interesting theorems about certain manifolds, those with properties P1, P2, P3, At the end of the student’s defence talk, a certain professor pointed out that the class of manifolds had no inhabitants, because some of the properties were mutually contradictory. According to the story, the professor phrased this question in an especially obnoxious way: “Isn’t your entire dissertation vacuous for the following trivial reasons?” If a judgment cannot be derived, it serves no purpose. However, a much more common problem in defining programming languages is that too many, or too few, judgments can be derived—not that *no* judgments of a given form can be derived.)

Forgetting about the horrible question, which is different from the other four because it asks for *all* possible instantiations of a meta-variable rather than just one, we find that a single judgment with two meta-variables e and v gives rise to $2 \times 2 = 4$ different problems, according to whether

- e is given or not given
- v is given or not given

Rephrasing the above questions (again, forgetting the horrible one):

1. The program that answers the question when e is given and v is not is an interpreter.
2. The program that answers the question when v is given, and e is not, doesn’t seem very useful, but we could call it a “reverse interpreter”.
3. The program that answers the question when both e and v are given is a program that *validates* whether a particular evaluation is correct.
4. The program that answers the question when neither e nor v is given is (hopefully) trivial, but it tells you that the judgment form isn’t vacuous (assuming at least one $e \Downarrow v$ judgment is in fact derivable).

§3 Defining a type system

3.5.2 Modes

If the instantiation of a meta-variable is given, we say its *mode* is *input*, and if it is not given, we say its mode is *output*. We can mark each meta-variable in a judgment form with the mode of that meta-variable. The “interpreter question” corresponds to the *moded judgment form*

$$e_{\text{IN}} \Downarrow v_{\text{OUT}}$$

and the “validation question” corresponds to

$$e_{\text{IN}} \Downarrow v_{\text{IN}}$$

(In Prolog and other logic programming languages, modes can be declared with $+$ and $-$; the IN mode corresponds to $+$, and the OUT mode corresponds to $-$).

For other judgment forms, we can also list various moded judgment forms. The small-step judgment form $e_1 \longrightarrow e_2$ is usually moded as

$$e_{1\text{IN}} \longrightarrow e_{2\text{OUT}}$$

but it’s reasonable to think about

$$e_{1\text{OUT}} \longrightarrow e_{2\text{IN}}$$

which corresponds to “expansions” in Church’s λ -calculus, whereas the $e_{1\text{IN}} \longrightarrow e_{2\text{OUT}}$ form corresponds to Church’s “reductions”.

For the typing judgment $\Gamma \vdash e : A$, the problem we’ve been considering (and playing with in Racket during lecture) corresponds to

$$\Gamma_{\text{IN}} \vdash e_{\text{IN}} : A_{\text{OUT}}$$

But other “modings” have been studied as well:

- The moding $\Gamma_{\text{IN}} \vdash e_{\text{IN}} : A_{\text{IN}}$ arguably is more appropriately called *type checking* than $\Gamma_{\text{IN}} \vdash e_{\text{IN}} : A_{\text{OUT}}$, which could be called *type inference*.
- An implementation of the moding $\Gamma_{\text{OUT}} \vdash e_{\text{IN}} : A_{\text{OUT}}$ would try to find a type *and* a context, which is both theoretically interesting and potentially useful: imagine a type error message that says, “I can’t type this, but *if only* this unknown identifier had a particular type...”. For example, you could ask such an implementation to come up with a context Γ and type A such that

$$\Gamma \vdash (\text{ite } (\text{id } x) (\text{id } y) (\text{id } z)) : A$$

and it might suggest $\Gamma = x : \text{Bool}, y : \text{Num}, z : \text{Num}$ and $A = \text{Num}$:

$$x : \text{Bool}, y : \text{Num}, z : \text{Num} \vdash (\text{ite } (\text{id } x) (\text{id } y) (\text{id } z)) : \text{Num}$$

which would be nice, especially when learning a language.

I know of one attempt to implement this, which I regard as a failure because it was extremely complicated, but perhaps there are simpler approaches.

- Finally, the moding $\Gamma_{\text{IN}} \vdash e_{\text{OUT}} : A_{\text{IN}}$ asks: under Γ , does *anything* have type A ? (In some type systems, you can write down a type that no expression has.¹)

¹However, in languages that have a type called `void`, `void` is usually *not* such a type, causing endless grumbling among academic researchers.

4 Declarative vs. algorithmic

The difficulty we’ve encountered with Type-lam (which we would encounter with Type-rec, as well) represents a gap between a rule that “looks good on paper” and one that can be directly implemented. Rules that *can* be directly implemented are called *algorithmic*: if you have enough practice, you can “read off” an algorithm from the rules. We have done this rather implicitly, but all of our implementations so far—of both evaluation and typing—have depended on being able to identify, without too much effort, which rule should be used. Usually only one rule has an expression of the right form in the conclusion: if the expression is (bfalse), only Eval-bfalse (for evaluation) or Type-bfalse (for typing) applies. Sometimes, two rules might apply, as with Eval-ite-true and Eval-ite-false; there, our interpreter uses the result of evaluating the scrutinee to decide which rule to use.

Until now (with Type-lam), we have not had to implement a rule in which a premise has missing information. Such rules are common in language definitions, or at least in the more theoretical work that underlies some languages (especially typed languages); type systems are connected to logics, but logicians are usually more interested in the theoretical properties of their rules than the connections to type systems. (This is historically true, at least, and logicians who predated the development of computers can hardly be faulted for not focusing on those connections!)

Rules that cannot be directly implemented, whether because of missing information or some other difficulty, are called *declarative*: they “declare” what the judgment means, but not “algorithmically”. A common tactic of programming languages researchers is to define a simple and (hopefully) understandable “declarative type system”, and *then* define an “algorithmic version” that *looks* very different but accepts exactly the same programs as the declarative type system.

We will not pursue that tactic now. Instead, we will change the definition of expressions in a way that provides the missing information in Type-lam (and in Type-rec). This isn’t my favourite solution, but it’s one that (I think) fits the time we have for Assignment 3; it’s also a solution that’s closer to what languages like C and Java do than the solution I like better.

5 Typing rules we *can* implement

The problem we have is that we don’t know what A is in Type-lam. So we’ll put A *into the expression*. The concrete syntax for lam will have a $\langle \text{Type} \rangle$, and the **define-type** branch will have an extra argument containing the domain of the function (the type of its argument); we can figure out the range of the function (the type of its result) by looking at the function body.

$$\frac{x : A, \Gamma \vdash e\text{Body} : B}{\Gamma \vdash (\text{lam } x \ \mathbf{A} \ e\text{Body}) : A \rightarrow B} \text{Type-lam} \qquad \frac{\Gamma \vdash e1 : A \rightarrow B \quad \Gamma \vdash e2 : A}{\Gamma \vdash (\text{app } e1 \ e2) : B} \text{Type-app}$$

We also have to do this for Type-rec, which makes me sad, because my preferred (but harder to explain) solution wouldn’t make us do this.

$$\frac{u : B, \Gamma \vdash e : B}{\Gamma \vdash (\text{rec } u \ \mathbf{B} \ e) : B} \text{Type-rec}$$

§5 Typing rules we *can* implement

Since a `rec` usually has a `lam` as its body, this means we have to write the function domain twice: once in the `rec`, as part of a function type, and once in the `lam`.

```
{rec u {-> Num Bool}
  {lam x Num
    {ite {= x 0}
      btrue
      {ite {= x 1}
        bfalse
        {ite {< x 0}
          {app u {+ x 2}}
          ; x must be >1
          {app u {- x 2}}}}}}}}
```